

- N-way replication doesn't improve read performance regardless of N
  - Might with unrealistically slow interconnect
- GFS moving from triplication to reed solomon directed from client
  - Should client handle replication or should server?
  - When client handles replication, more possibility to introduce jitter
  - Perhaps jitter will be better tolerated at exascale
  - Replication must be coordinated. Should coordination attempt to optimize space or optimize bandwidth?
- In GFS, clients and servers are the same. Not true for HPC.
- Workload dependent
  - Are different replication schemes more appropriate for different workloads?
  - What are assumptions behind a particular scheme?
  - Do MR/HDFS schemes favor read-intensive whereas HPC might be more write-intensive?
- What about co-design?
  - facebook has multiple clusters for different purposes: one for mem cache, one for Hadoop, one for MySQL
  - HPC has three missions: nuclear security, nuclear energy, climate change
    - At first blush, homogenous enough to require just one design
    - Should this be re-examined?
  - Dangerous for shifting workloads